

Privacy-Protective-GAN for Privacy Preserving Face De-Identification

Yifan Wu¹, Fan Yang¹, Yong Xu², *Senior Member, CCF, ACM, IEEE*, and Haibin Ling¹, *Senior Member, IEEE*

¹*Department of Computer and Information Sciences, Temple University, Philadelphia 19122, U.S.A.*

²*School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China*

E-mail: {yifan.wu, fyang}@temple.edu; yxu@scut.edu.cn; hbling@temple.edu

Received July 12, 2018; revised December 20, 2018.

Abstract Face de-identification has become increasingly important as the image sources are explosively growing and easily accessible. The advance of new face recognition techniques also arises people's concern regarding the privacy leakage. The mainstream pipelines of face de-identification are mostly based on the k -same framework, which bears critiques of low effectiveness and poor visual quality. In this paper, we propose a new framework called Privacy-Protective-GAN (PP-GAN) that adapts GAN (generative adversarial network) with novel verifier and regulator modules specially designed for the face de-identification problem to ensure generating de-identified output with retained structure similarity according to a single input. We evaluate the proposed approach in terms of privacy protection, utility preservation, and structure similarity. Our approach not only outperforms existing face de-identification techniques but also provides a practical framework of adapting GAN with priors of domain knowledge.

Keywords face de-identification, privacy protection, image synthesis, generative adversarial network (GAN)

1 Introduction

Beneficial from the blooming development of media and network techniques that make a vast amount of images more approachable, image analysis techniques bear their prosperity in the past decade and bring unprecedented convenience to our daily life. Among those image sources exposed to the public with or without our awareness, a considerable number of them contain our identity especially the biometric information. Not only will the unprotected exposing cause the leak of privacy, the common approaches of protection like blurring and pixelization may also not be satisfied in thwarting face recognition software^[1,2]. The other extreme side is that we simply mask identity area off, which is perfect for identity removal. But it causes serious loss of data utility in application of visual understanding as the scene information is changed with objects removal. Thus it is critically important to build a framework that can properly de-identify the privacy information from the image while keeping its utility at the same time.

Specially for the face de-identification problem, the dilemma is that on the one hand, the de-identified images are required to look different from the original image to ensure the removal of identity; on the other hand, we expect the de-identified image to retain as much structural information in the original image as possible so that the image utility remains. Previous approaches on this problem are mostly based on k -same algorithm^[1,3–5] for the de-identification procedure and some apply additional models like the Active Appearance Models (AAMs)^[6] to explicitly construct faces preserving the utility attributes like gender, race, and age^[7,8]. Yet these models fail to make full use of existing data and deliver fairly poor visual quality due to the unnatural synthesis.

The generative adversarial networks (GANs) provide an inspiring framework on generating sharp and realistic natural image samples via adversarial training^[9–11]. GAN can be naturally used for the face de-identification as it can generate new samples from the gallery following original input data distribu-

tion. Since GAN implicitly models a dataset distribution, it cannot involve processing individual images. To achieve an image-to-image face de-identification, conditional GAN (cGAN) is intuitively leveraged here, which is an extension of GAN that fits a conditional data distribution.

Unfortunately, cGAN is not directly applicable to our scenario, which requires privacy protection as well as data utility preservation. First, the generative loss in cGAN is not specific for distinguishing identities; thus the generated samples could be either too similar to each other to pass external face verification, or visually ghosted as none-faces. Second, the structure information is unensured to be preserved due to the inconsistency in feature spaces between GAN-type and structural losses.

Thus to make an adaptive GAN framework for the de-identification problem, we introduce additional mechanisms to enlarge the sampling range away from the given samples in the identity-related embedding feature space properly and to constrain the variation of the part of image structure from the full image feature space simultaneously. When unified in the GAN loss, these two tasks are actually contradictory to each other because the GAN loss is unable to explicitly tell the embedding feature from the structure feature and they contribute to the performance of distinguishing two faces in the same direction. Thus here we explicitly express them as external losses to compensate their contradiction in GAN and consequently achieve a novel Privacy-Protective-GAN model (PP-GAN, see Section 3), which can plausibly balance between id-removal and structure retain. Fig.1 represents the original images and corresponding de-identified images.

Indeed, in our PP-GAN model, we introduce two additional types of external modules: the verifier with contrastive loss^[12] to allow cGAN's generator for a larger range of sampling exploration, and the regulator with Structural Similarity Index (SSIM)^[13] to account for the structure preservation. Intuitively, the verifier adds the prior information of identity feature matching in the embedding space to help remove biometric information while the regulator tells the generator how to maintain similar image utility via luminance, contrast and structural similarities. The involvement of these two types of prior knowledge on what the identity information versus the structure information is significantly improves the model performance.

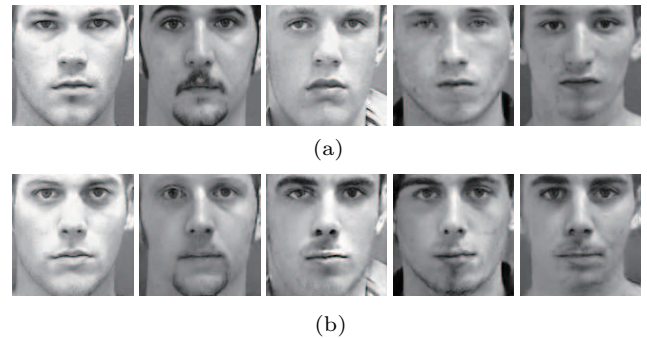


Fig.1. Illustration of face de-identification. (a) Original face images. (b) Corresponding de-identified images synthesized by our method.

In the experimental setting, we quantitatively demonstrate the effectiveness of proposed method in several aspects: 1) the generative de-identified image can thwart the face verifier; 2) the de-identified image is not simply switched with others in the training data; 3) the luminance, contrast and structure can be partially preserved; 4) the attributes utility can be preserved when training an attribute specific generator.

In summary, our contributions are as following.

- To the best of our knowledge, we are the first to propose a GAN-based framework that is trainable in an end-to-end manner directly for the face de-identification task. The integrated framework can synthesize the de-identified output for each single probe face.

- We manage to balance the trade-off between image quality and privacy protection by introducing novel modules, the verifier, and the regulator. The model learns to retain the structure similarity with the original image on the pix level, while distinguishing the input and synthesized output in the identity-related feature space. It is the first time that a de-identification metric is properly aggregated into the objective function in a controllable and measurable manner, which largely ensures the effectiveness of de-identification and privacy protection.

- We are the first to propose a systematic plan of evaluating the results of de-identification, where we employ a state-of-the-art face verifier^[14] to determine the de-identification rate and then a face detector^[15] to quantify the detection rate. These two measurements cooperatively indicate how well the de-identification procedure works in its target with utility retained. The proposed method achieves promising de-identification rate and obtains the best trade-off between identity removal and visual similarity.

The rest of the paper is organized as follows. Related work about de-identification task is summarized

in Section 2. In Section 3, we introduce our PP-GAN framework. A quantitative evaluation on a public dataset as well as visual illustrations is provided in Section 4. The paper concludes with a discussion in Section 5.

2 Related Work

Privacy protection in visual data has been drawing increasing amount of research attention recently, with focuses on different scenarios such as video analysis and security and sampled studies in [16-22].

Face de-identification is an important tool for visual privacy protection. The goal of face de-identification has two folds: privacy protection and data utility preservation. In such way, de-identified images conceal the identity privacy of the original image, while preserving non-identity-related aspects for further data analysis. Earlier studies on face de-identification simply use masking, blurring, or pixelation^[23]. While these Ad Hoc methods are easily applicable, it is shown that they provide no privacy assurance since they may fail to thwart face recognition software^[24]. Moreover, how to conduct a sufficient blur itself is non-trivial^[25].

To address this problem, Newton *et al.*^[1] proposed the k -same algorithm based on the k -anonymity concept^[26]. By applying the k -same algorithm, a given image is represented by average face of k -closest faces from the gallery. This procedure theoretically limits the performance of recognition to $1/k$, but usually suffers from ghosting artifacts in de-identified images. Some variants of k -same are proposed to improve the data utility and the naturalness of de-identified face images. The k -Same-Select algorithm^[3] partitions the image set into mutually exclusive subsets and applies the k -same algorithm independently to each subset, aiming to preserve the attributes of each subset. In order to overcome undesirable artifacts due to misalignments and produce de-identified images of better quality, the k -Same-M algorithm^[4] based on active appearance models (AAMs)^[6] is proposed. This algorithm fits an AAM to input images and then applies k -same to the AAM model parameters. To further keep the data practically useful, Du *et al.* explicitly preserved race, gender and age attributes in face de-identification^[7]. Jourabloo *et al.* proposed to jointly model face de-identification and attribute preservation in a unified optimization framework^[8].

The majority of state-of-the-art approaches are designed based on the k -same and implemented using

AAM models. However, these methods have notable limitations. 1) The k -same assumes that each subject is only represented once in the dataset, but this may be violated in practice. The presence of multiple images from same subject or images sharing similar biometric characteristics can lead to lower levels of privacy protection. 2) The k -same operates on a closed set and produces a corresponding de-identified set, which is not applicable in situations that involve processing individual images or sequences of images. 3) Generated images by AAM models do not yet look natural enough.

Recently, GANs have presented promising effectiveness on image generating problems. It is a natural tool to synthesize de-identified images. Brkic *et al.*^[27] built a GAN-based model to generate full body images for de-identification, but the quality in face areas is not guaranteed. Blaž *et al.* used GAN to synthesize de-identified faces, but still based on the k -same algorithm^[22,28,29]. [22] leverages an auxiliary loss to confound the gender attribute while retaining the biometric recognition performance.

In this paper, we only focus on biometric information, i.e., we do not remove soft biometric features like age or race and non-biometric information like hairstyles^[2]. We propose a novel Privacy-Protective-GAN framework to address the problems mentioned above.

3 Method

Face de-identification can be formulated as a transformation function δ which maps a given face image x to a de-identified image \hat{x} , i.e., $\delta(x) = \hat{x}$, aiming to mislead the face verification. In this work, we propose a PP-GAN model and adopt the final generator as the de-identification function. First, we pretrain a verifier to determine whether two faces are from the same subject. Next, we freeze this verifier and utilize it to compute the contrastive loss in the whole system training phase. In each iteration, the synthesized output and the original image will be forwarded in: 1) the identity-related feature space by the pre-trained verifier to enforce the removal of identity; 2) the pixel level structure similarity by SSIM to preserve visual correspondence. Then the contrastive loss and the SSIM loss will be backpropagated to update the generator. Fig.2 gives the big picture of the whole framework.

In the following subsections, we will describe the structure of PP-GAN and explain its learning procedure in detail. The overview of the whole system is illustrated in Subsection 3.1. The detailed explanations

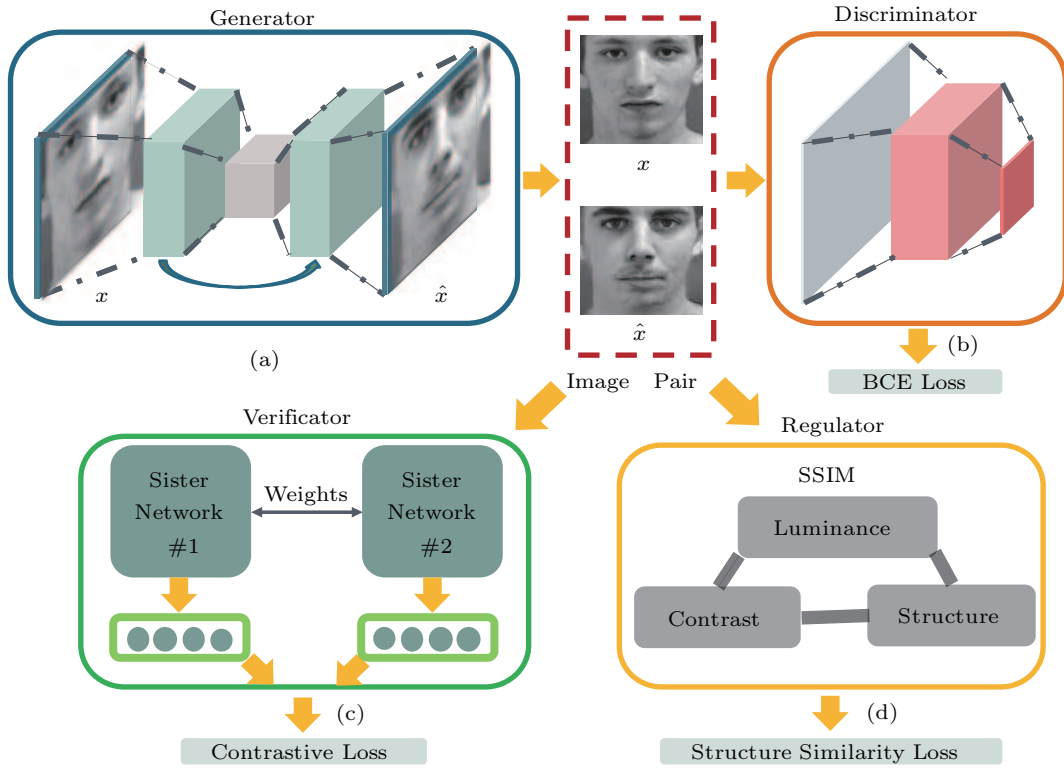


Fig.2. Structure of the privacy-protective-GAN (PP-GAN) framework. The PP-GAN framework consists of four components: generator, discriminator, verifier, and regulator. (a) The generator is a “U-Net” based auto-encoder, synthesizing a de-identified image \hat{x} for given original image x . (b) The discriminator is adversarially trained with the generator to encourage the output to be sharp and realistic. (c) The verifier is to determine whether two faces are from the same person. (d) The regulator is to determine how similar the two faces are in pixel level.

of the generator vs the discriminator are given in Subsection 3.2. The verifier and the structure similarity are introduced in Subsections 3.3 and Subsections 3.4 respectively.

3.1 Objective

Face de-identification is a complex task involving multiple constraints at the same time. Specially, given an original face, we generate a new face image accordingly and hope that it meets three types of qualification. 1) The generated face should look natural and realistic, i.e., it looks like a face. 2) The generated face should not be discriminated as the same person with the original face, i.e., it should be far enough from the original one in the identity-related feature space. 3) The de-identified image should be consistent with the original image on the pixel level as much as possible, i.e., it shares similar luminance, contrast, and structure with the original image.

For the first objective, it can be achieved by utilizing the cGAN, which adversarially trains a generator

(G) taking the original image as input and producing de-identified images versus the co-trained discriminator (D). The cGAN loss is denoted as $\mathcal{L}_{cGAN}(G, D)$ here. The details of cGAN settings will be clarified in Subsection 3.2.

For the second goal, we pre-train a verifier to produce a verification loss in the identity embedding space for each pair of original and de-identified images. The verification loss for each generator G is then formalized as the expected loss across subjects and denoted as $\mathcal{L}_{verif}(G)$. This is used to guide the generator to enforce that the output holds an enough distance with the input in the identity-related feature space. The details about the verifier will be elaborated in Subsection 3.3.

For the third objective, we adopt the structural similarity index (SSIM)^[13] to quantify the similarity between the original and the generated images. The SSIM loss combines the luminance, contrast, and structural differences in a product form with different power coefficients. Similarly, as the verification loss, the SSIM

loss for the generator G is then formularized as the expected loss across subjects and denoted as $\mathcal{L}_{\text{sim}}(G)$. The details of structure loss will be discussed in Sub-section 3.4.

Combining these three types of losses, our final objective of face de-identification is defined as

$$\begin{aligned} &\mathcal{L}_{\text{face}}(G, D) \\ &= \mathcal{L}_{\text{cGAN}}(G, D) + \lambda_1 \mathcal{L}_{\text{verif}}(G) + \lambda_2 \mathcal{L}_{\text{sim}}(G), \end{aligned} \quad (1)$$

and the optimal generator G^* is solved through the min-max procedure

$$G^* = \arg \min_G \max_D \mathcal{L}_{\text{face}}(G, D),$$

where λ_1 and λ_2 are the hyperparameters for multiple losses.

3.2 Architecture of GAN

The cGAN learns a mapping $G : \{x, z\} \rightarrow \hat{x}$ from an observed image x with additional random noise z to a synthesized image \hat{x} , where x is referred to as a “real” sample or the condition from the original dataset, \hat{x} is referred to as a “fake” sample generated by the trained generator G , and z is the random noise to ensure the image variability. The adversarial procedure trains a generator to produce outputs that can hardly be distinguished as a “fake” by the co-trained discriminator (D). Mathematically, the objective function of the cGAN is expressed as

$$\begin{aligned} &\mathcal{L}_{\text{cGAN}}(G, D) \\ &= \mathbb{E}_{x, \hat{x} \sim P_{\text{data}}(x, \hat{x})} (\log D(x, \hat{x})) + \\ &\quad \mathbb{E}_{x \sim P_{\text{data}}(x), z \sim P_z(z)} (\log(1 - D(x, G(x, z)))) \end{aligned}$$

where G is the generator and D is the discriminator. The optimal G^* minimizes this objective against an adversarial D that maximizes it, and it can be solved via a min-max procedure $G^* = \arg \min_G \max_D \mathcal{L}_{\text{cGAN}}(G, D)$.

In practice, we adapt our generator and discriminator architectures from the Image-to-Image translation^[10]. As shown in Fig.3, both of them use modules consisting of Convolution-BatchNorm-ReLu^[30]. The generator is a “U-Net”^[31] by adding skip connections between symmetric layers in the convolution and deconvolution steps. The stride keeps to be 2; thus it resizes by 2 in the encoding part, and expands by 2 in the decoding part. The size of the embedding layer is 1×1 . Two drop-out layers in the middle serve as random noise z . For the discriminator, instead of feeding it with “real” or “fake” pairs, we follow the standard approach from [9] and input either “real” image x or “fake” image \hat{x} since there is no unique solution for the de-identified $G(x)$. Also, a patch-design^[10] is adopted in the discriminator; thus it outputs an $N \times N$ patch, rather than a single value to represent the probability of current input to be “real”. Here we set N to 30 with receptive field of size 34. The network architectures are shown in Fig.3.

3.3 Face Verification

Face verification is to determine whether two images express the same person. Metric learning is a widely used approach for this task, which learns semantic distance measurements and the associated embeddings^[32–35] from the original image space to feature space. Motivated by FaceNet^[34], we strive for an embedding $f(x)$ so that the squared distance among

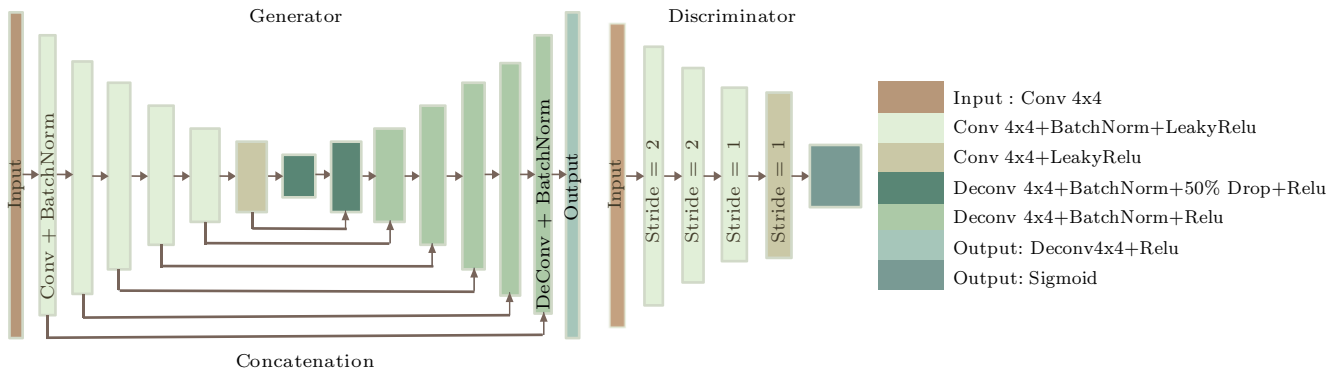


Fig.3. Network architectures of the generator G and the discriminator D . The generator has a “U-net” architecture with the input size of $1 \times 128 \times 128$, the embedding size of $256 \times 1 \times 1$, and the output size that is consistent with the input size. The input size of the discriminator is $1 \times 128 \times 128$, while the output size is $1 \times 30 \times 30$.

all faces is independent of imaging conditions. In the embedding spaces, samples with the same identity are closer whereas those with different identities are further apart.

Mathematically, the embedding is denoted as f : $\text{data} \rightarrow \mathbb{R}^d$, which embeds an image x into a d -dimensional Euclidean space. We adopt the following contrastive loss function originally proposed by Hadsell *et al.*^[12] The loss function is defined as

$$\begin{aligned} & \text{Verif}(x_i, x_j, \eta_{x_i, x_j}, \alpha) \\ = & \begin{cases} \frac{1}{2} \|f(x_i) - f(x_j)\|_2^2, & \text{if } \eta_{x_i, x_j} = 0, \\ \frac{1}{2} \max(0, \alpha - \|f(x_i) - f(x_j)\|_2)^2, & \text{if } \eta_{x_i, x_j} = 1, \end{cases} \end{aligned} \quad (2)$$

where α is the margin that is enforced between positive and negative pairs, f is the embedding function, and η_{x_i, x_j} is the indicator where $\eta_{x_i, x_j} = 0$ means positive pairs, and $\eta_{x_i, x_j} = 1$ means negative pairs.

For the network architecture, we choose the new Light CNN-9 model^[36], which introduces the MaxFeature-Map (MFM) operation as an alternative of ReLU and claims that MFM can play a role of feature selection. With this design, the Light CNN-9 model achieves state-of-the-art results with less parameters and time-consumption. The Light CNN-9 model contains five convolution layers, four Network-in-Network (NIN) layers^[37], Max-Feature-Map function, and four max-pooling layers. The final layer is a fully connected one that outputs a 256-dimensional representation. A detailed explanation of this architecture can be found in [36].

In the pre-training phase of the verifier, we randomly sample 50% positive and 50% negative pairs as train data to avoid data unbalance. These pairs are separately fed to a shared weight Siamese network, where the model is updated by the contrastive loss defined in (2). Once the verifier training is finished, parameters are frozen in the whole system training followed.

Next in the PP-GAN training phase, the pair of original image x and synthesized image \hat{x} is fed into pre-trained verifier with $\eta_{x, \hat{x}} = 1$ to enforce their different identities, i.e., the de-identification.

The term of verification loss in (1) is then defined as

$$\mathcal{L}_{\text{verif}}(G) = \mathbb{E}_{x \sim P_{\text{data}(x)}, z \sim P_z} (\text{Verif}(x, G(x, z), \eta, \alpha)),$$

where x is the original input I , z is the noise, $G(x, z)$ is the generated image, $\eta = 1$ is the always false indicator, and the margin is set to 2 on the normalized feature vector to enforce a positive loss.

3.4 Structure Similarity

For the de-identification problem, we want to hold the correspondence between original image x and synthesized image \hat{x} . The vanilla GAN simply generates new samples from the data distribution. For cGAN, although it samples under the condition of the input image x , it does not ensure that the synthesized image \hat{x} is close enough to x in the image space. For example, if we have two images x_i and x_j with the same condition, the cGAN would generate \hat{x}_i and \hat{x}_j with the same distribution so that we cannot tell whether \hat{x}_i is the de-identified image from x_i or x_j . Thus an additional structural loss is necessary here to regulate the generator to assure this matching. The two popular choices, the mean squared error (MSE) and the related quantity of peak signal-to-noise ratio (PSNR), are not well suited because they do not match very well to the perceived visual quality by humans, and the pixel-level matching with the reference image is not what we want to optimize. Instead we use the structural similarity index (SSIM)^[13], which has been used previously for privacy protection^[38], as an objective measurement for assessing perceptual image degradation after de-identification.

SSIM consists of three components including luminance similarity $l(x, \hat{x})$, contrast similarity $c(x, \hat{x})$, and structural similarity $s(x, \hat{x})$:

$$\text{SSIM}(x, \hat{x}) = l(x, \hat{x})^\alpha \times c(x, \hat{x})^\beta \times s(x, \hat{x})^\gamma.$$

For a detailed explanation of SSIM, one can refer to [13]. Here we define $\mathcal{L}_{\text{sim}}(x, \hat{x}) = \frac{1}{2}(1 - \text{SSIM}(x, \hat{x}))$. Then in the GAN training, we have

$$\mathcal{L}_{\text{sim}}(G) = \mathbb{E}_{x \sim P_{\text{data}(x)}, z \sim P_z} \left(\frac{1}{2} (1 - \text{SSIM}(x, G(x, z))) \right),$$

where z is the noise, and $G(x, z)$ is the generated image.

4 Experiments

4.1 Experimental Setup

We evaluate the proposed framework on a publicly available MORPH^[39] dataset, which contains 55 000 unique images of more than 13 000 individuals, with diverse demographic information like age, gender, and

race. Here we only use male data since the number of female subjects is limited. The details of the demographic distribution of the MORPH dataset are shown in Fig.4.




Fig.4. Demographic distribution of the MORPH dataset. The black group contains more samples than the white group, while both have the similar age distribution.

To validate that our framework can preserve attributes utility, data is divided into black and white according to race, and then we divide each category into three age groups: the youth (age ≤ 25), the middle-aged ($25 < \text{age} < 40$) and the senior (age ≥ 40). We regard these eight different data groups as black, white, black-youth, black-middle, black-senior, white-youth, white-middle, and white-senior. We randomly pick 90% from each group for training and use the remaining for testing. The training and the testing sets are consistent in the verifier pretraining and the whole system training. For the data preprocessing, we use OpenFace^① to detect, align, and crop face areas. During verifier pretraining, we set initial learning rate to 1×10^{-4} and margin to 2. During the whole system training, the learning rate is set to be 1×10^{-5} . For all the experiments, we use minibatch SGD and apply the Adam solver^④.

For each group, we separately conduct four experiments to analyze our objective function: cGAN, cGAN with SSIM loss, cGAN with pretrained verifier, and cGAN with both pretrained verifier and SSIM constraint (i.e., ours). Here we regard these different setups as: cGAN-only, cGAN+Sim, cGAN+Verif, and cGAN+Sim+Verif.

4.2 Privacy Protection

To demonstrate the proposed framework that successfully generates face images for privacy protection, we need to prove that the produced face images neither remain the same as the corresponding raw images,

which means successful de-identification, nor simply change the identity to other identities in the dataset, which means no switching identity occurs.

De-Identification. To evaluate if the proposed model successfully removes the identity of original images, we assess if a pair of images are from the same identity in a verification way. Specifically, given a verifier, for each pair of images (x, \hat{x}) , it gives a distance value indicating the similarity of the two images. If the value is smaller than a threshold, the two images are treated as the same identity, which means de-identification fails; otherwise, the two images are determined as different identities, and then de-identification succeeds. Here we use de-identification rate (ERR rate) to represent the effectiveness of privacy protection, i.e., (x, \hat{x}) is not determined as the same person.

To eliminate the overfitting caused by the pre-trained verifier since it guides the training of the whole system, we fine-tune the FaceNet^② with triplet loss as our evaluation verifier. FaceNet is trained on the eight subgroups respectively. Specifically, the test data is grouped into numerous pairs with half positive and half negative. Then the trained model is evaluated on these pairs of images in a 10-fold way. The finally gained optimal threshold is utilized in de-identification evaluation.

Table 1 and Table 2 show our de-identification performance. FaceNet achieves verification accuracy greater than 97% on all black groups and 94% on white groups. The overall de-identification rates are both very high, and that of the black group is a bit higher due to the larger sample size.

Table 1. De-Identification Rate (%) of Male/Black

	Original	De-Identified	De-Identified
	Test	Train	Test
Black	1.5	100.0	100.0
Black-Youth	1.6	100.0	100.0
Black-Middle	1.6	100.0	97.0
Black-Senior	2.7	97.2	93.7

Table 2. De-Identification Rate (%) of Male/White

	Original	De-Identified	De-Identified
	Test	Train	Test
White	1.8	100.0	100.0
White-Youth	2.8	96.2	94.4
White-Middle	3.3	89.7	90.8
White-Senior	5.9	91.7	84.7

^①<https://cmusatyalab.github.io/openface>, Dec. 2018.

^②<https://github.com/davidsandberg/facenet>, Dec. 2018.

No Switching-Identification. In addition, in order to demonstrate the yielded image does not possess the identity of the other original images, for each generated image, we compute its similarity, i.e., distance, with every original image, and compare the distance of image pair (x_i, \hat{x}_j) with a threshold, where i, j are the identity of the image. If the distance is smaller than the threshold and $i \neq j$, identification switching occurs. If the distance is smaller than the threshold and $i = j$, identification switching does not occur. We use the amount of identification switches (IDS) as a measurement.

We conduct experiments on black-youth and white-youth. As the data is unbalanced, to fairly compare the IDS on both groups, we randomly sample the same number of people from black-youth test as that in the white-youth test. On both groups, we obtain 0 IDS, which means no identification switching occurs.

4.3 Utility Preservation

Face-Like. The most important point for data utility is that the generated images should look natural and realistic, i.e., they look like face. Here we utilize a face detector MTCNN^[15] to determine whether the generated images are face or not. To alleviate the variance of the face detector, we conduct experiments on both original and generated face images. Table 3 and Table 4 show the detection rate on both the original and the generated data. From the tables, we note that the detection rates on original and de-identified data are similar, proving that our system preserves the face-like utility very well. As cropped face images are not suitable for face detector trained on face image with context, to alleviate the effects from the inconsistency, we pad the cropped image with 50 pixels of 0 along each axis. For some groups, the detection rate on de-identified images is higher than that on original images, e.g., black-youth. This is because some original face images are distorted, but yielded images have less distorted images.

We show some representative bad cases in Fig.5. The proposed method might fail when meeting images with a non-frontal head pose or with a large beard (occlusion) as the training data is normalized, and mostly frontal and without occlusion.

Attributes Preservation. To examine the performance of the proposed model for attribute preservation, we train separate classifiers for each attribute and compute the classification accuracy for the generated images. The classification accuracy rate is used as a measurement to evaluate experimental results. From Fig.6,

we can see that the age-specific models well preserve the group attributes compared with the age-nonspecific group. The experiments on original images display similar comparisons.

Table 3. Detection Rate of Male/Black

Data Setting	Black	Black-Youth	Black-Middle	Black-Senior
Original w/o padding	0.733	0.817	0.767	0.598
De-identified w/o padding	0.719	0.847	0.715	0.570
Original w padding	0.988	0.993	0.994	0.977
De-identified w padding	0.956	0.995	0.990	0.969

Note: w: with; w/o: without.

Table 4. Detection Rate of Male/White

Data Setting	White	White-Youth	White-Middle	White-Senior
Original w/o padding	0.758	0.808	0.738	0.621
De-identified w/o padding	0.784	0.767	0.739	0.499
Original w padding	0.973	0.967	0.979	0.975
De-identified w padding	0.975	0.940	0.988	0.902



Fig.5. Failed examples. We present some representative failed examples here.

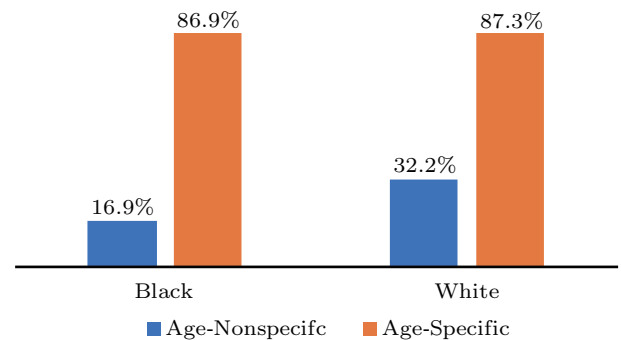


Fig.6. Age classification accuracy. The blue bar displays the classification accuracy of preserving attributes for age-nonspecific group and the orange bar is for age-specific group.

4.4 Visual Similarity

For the de-identification problem, in addition to the utility preservation, we want to hold the correspondence between original image x and synthesized image \hat{x} rather than just replace the face area randomly. In other words, we only want to remove the privacy related characteristics, but keep the visual similarity, e.g., contours and luminous condition, as much as possible.

As shown in Fig.7, the original and the generated images share consistent SSIM-consisting features like luminance, contrast, and structure information while distinguishing privacy-related characteristics such as the sizes and shapes of eyes, nose, and mouth. This indicates the effectiveness of the regulator.




Fig.7. Visualization of correspondence between original and de-identified images. The upper row shows the original face images, and the lower row shows the corresponding de-identified images synthesized by our PP-GAN model. (a) ID: 325273. (b) ID: 090616.

To demonstrate that the proposed framework achieves both de-identification and visual similarity, we compute id-removal rate and SSIM of each yielded image with the corresponding original image on cGAN-only, cGAN+Verif, cGAN+Sim, and cGAN+Verif+Sim (see Fig.8), and display representative results (see Fig.9). Overall, our framework maintains high id-removal rate and SSIM at the same time.

One can also visually tell that the PP-GAN model maps the original face to a different face while somewhat keeping certain implicit relationship.

For the cGAN-only setting, it shows a high ID-removal rate with uncertain SSIM for the black (Fig.8(a)) and a low id-removal rate with high SSIM for the white (Fig.8(b)), which may come from different sizes of samples. When the sample size is small, the cGAN-only generates new images quite similar to existing one (Fig.9(a)) and gains high SSIM. When the sample size is large, the cGAN-only generator is able to step away from the original figures but may generate improper visualizations (Fig.9(b)), and provides low insurance for SSIM.

For the cGAN+Sim setting, the SSIM cost degenerates cGAN to the same manner of cGAN with limited samples, which functions almost as an identity mapping (Fig.9). Thus it holds low id-removal rate and high SSIM.

For the cGAN+Verif setting, the verification loss here enlarges the cGAN’s sampling range and thus results in a high id-removal rate (Fig.8), while the inconsistency between the two feature spaces and the lack of structure constraints may cause the visually strange consequences (Fig.9(b)). An interesting phenomenon is that when the sample size is small (the white group),

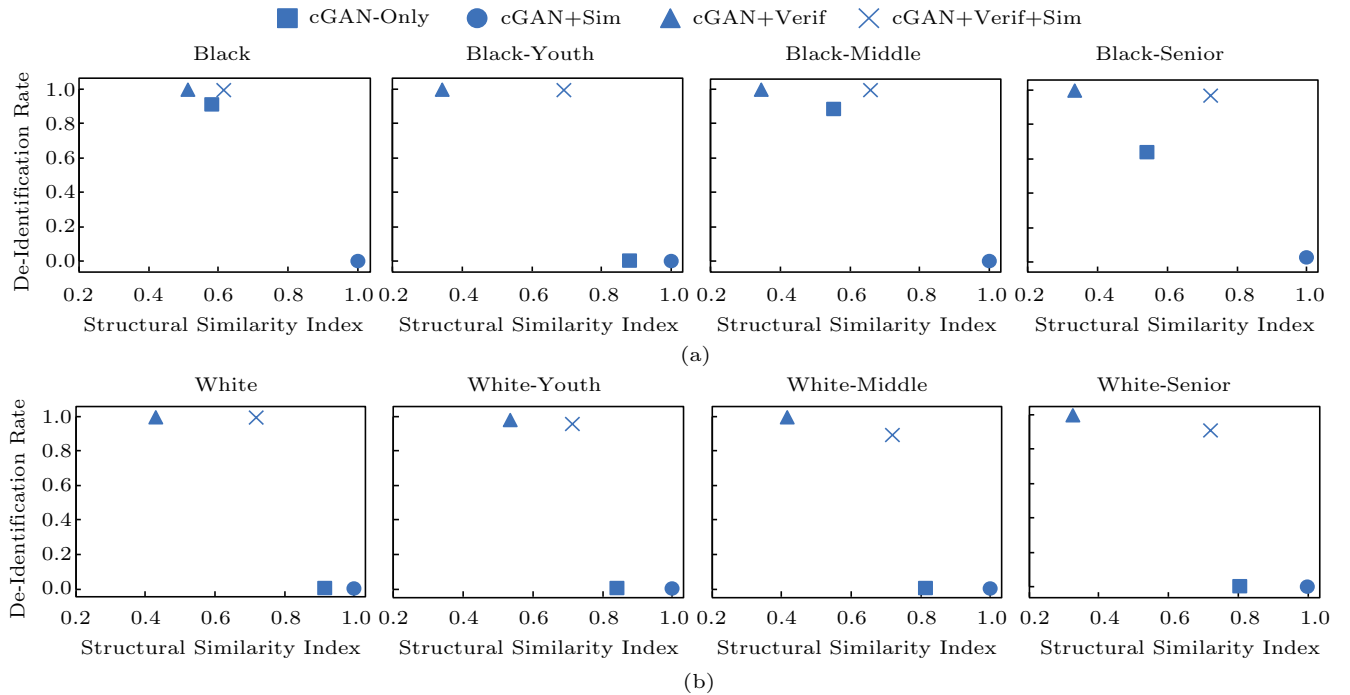



Fig.8. Trade-off between identity removal and visual similarity. The horizontal axis represents values of structural similarity index (SSIM), and the vertical axis represents the de-identification rate.



(a)



(b)

Fig.9. Visualized results of different objective designs on different groups. Our PP-GAN generates visually agreeable results beyond all the other settings on every group.

the cGAN+Verif indeed obtains a visually moderate result (Fig.9(a)) due to the enlarged sampling range from the verifier.

These results confirm the necessity and efficiency of simultaneously including the verifier and regulator to balance the id-removal and utility retaining.

5 Conclusions

In this paper, we presented a new face de-identification framework, Privacy-Protective-GAN,

which generates de-identified output according to a single input. We explicitly integrated the de-identification metric into the objective function to ensure the privacy protection. Meanwhile, we tried to preserve visual similarity as much as possible to retain data utility by adding a regulator. In the experiments, we quantitatively demonstrated the effectiveness of the proposed method in terms of privacy protection, utility preservation, and visual similarity. More qualitative results are shown in Fig.10 and Fig.11.



(a)



(b)



(c)



(d)

Fig.10. Visualized results of the black group. In each pair of double rows, the top one is for the original faces and the bottom one is for the de-identified faces. (a) This panel shows the results for the whole black group. We can see that the age factor is not retained. (b) shows the results of black-youth subgroup. (c) shows the results of black-middle subgroup. (d) shows the results of black-senior subgroup.



(a)



(b)



(c)



(d)

Fig.11. Visualized results of the white group. In each pair of double rows, the top one is for the original faces and the bottom one is for the de-identified faces. (a) This panel shows the results for the whole white group. We can see that the age factor is not retained. (b) shows the results of white-youth subgroup. (c) shows the results of white-middle subgroup. (d) shows the results of white-senior subgroup.

Limitations and Future Work. Although our approach generates realistic de-identified output, and achieves promising de-identification rates, it struggles in applying on faces in the wild with head poses, occlusions and so on as training data is mostly frontal. Furthermore, we present a potential to preserve soft-biometric attributes by training models on each attributes-specific subgroup. It is reasonable to unify the model and control attributes selection in the future exploration. Finally, extending this work to videos

data, achieving temporal consistency would be an interesting direction.

References

- [1] Newton E M, Sweeney L, Malin B. Preserving privacy by de-identifying face images. *IEEE Transactions on Knowledge and Data Engineering*, 2005, 17(2): 232-243.
- [2] Ribaric S, Ariyaeinia A, Pavesic N. De-identification for privacy protection in multimedia content: A survey. *Signal Processing: Image Communication*, 2016, 47: 131-151.

- [3] Gross R, Airoldi E, Malin B, Sweeney L. Integrating utility into face de-identification. In *Proc. the 5th International Conference on Privacy Enhancing Technologies*, May 2005, pp.227-242.
- [4] Gross R, Sweeney L, de la Torre F, Baker S. Model-based face de-identification. In *Proc. the 2006 Conference on Computer Vision and Pattern Recognition*, Jun. 2006, Article No. 161.
- [5] Gross R, Sweeney L, de La Torre F, Baker S. Semi-supervised learning of multi-factor models for face de-identification. In *Proc. the 2008 Conference on Computer Vision and Pattern Recognition*, Jun. 2008, Article No. 10.
- [6] Matthews I, Baker S. Active appearance models revisited. *International Journal of Computer Vision*, 2004, 60(2): 135-164.
- [7] Du L, Yi M, Blasch E, Ling H. GARP-face: Balancing privacy protection and utility preservation in face de-identification. In *Proc. IEEE International Joint Conference on Biometrics*, Sept. 2014, Article No. 67.
- [8] Jourabloo A, Yin X, Liu X. Attribute preserved face de-identification. In *Proc. IEEE International Conference on Biometrics*, May 2015, pp.278-285.
- [9] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y. Generative adversarial nets. In *Proc. Annual Conference on Neural Information Processing Systems*, Dec. 2014, pp.2672-2680.
- [10] Isola P, Zhu J Y, Zhou T, Efros A A. Image-to-image translation with conditional adversarial networks. In *Proc. the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, Jul. 2017, pp.5967-5976.
- [11] Zhu J Y, Park T, Isola P, Efros A A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, Jul. 2017, pp.2242-2251.
- [12] Hadsell R, Chopra S, LeCun Y. Dimensionality reduction by learning an invariant mapping. In *Proc. the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Jun. 2006, pp.1735-1742.
- [13] Wang Z, Bovik A C, Sheikh H R, Simoncelli E P. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004, 13(4): 600-612.
- [14] Schroff F, Kalenichenko D, Philbin J. FaceNet: A unified embedding for face recognition and clustering. In *Proc. the 2015 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2015, pp.815-823.
- [15] Zhang K, Zhang Z, Li Z, Qiao Y. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 2016, 23(10): 1499-1503.
- [16] Sun Q, Tewari A, Xu W, Fritz M, Theobalt C, Schiele B. A hybrid model for identity obfuscation by face replacement. In *Proc. the 15th European Conference on Computer Vision*, Sept. 2018, pp.570-586.
- [17] Chen D, Chang Y, Yan R, Yang J. Tools for protecting the privacy of specific individuals in video. *EURASIP Journal on Advances in Signal Processing*, 2007, 2007(1): Article No. 075427.
- [18] Wilber M J, Boulton T E. Secure remote matching with privacy: Scrambled support vector vaulted verification (S^2V^3). In *Proc. IEEE Workshop on Applications of Computer Vision*, Jan. 2012, pp.169-176.
- [19] Chan A B, Liang Z S, Vasconcelos N. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *Proc. the 2008 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2008, Article No. 229.
- [20] Sun Q, Ma L, Oh S J, van Gool L, Schiele B, Fritz M. Natural and effective obfuscation by head inpainting. In *Proc. the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun. 2018, pp.5050-5059.
- [21] Sun Q, Schiele B, Fritz M. A domain based approach to social relation recognition. In *Proc. the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, Jul. 2017, pp.435-444.
- [22] Mirjalili V, Raschka S, Namboodiri A, Ross A. Semi-adversarial networks: Convolutional autoencoders for imparting privacy to face images. In *Proc. the 2018 International Conference on Biometrics*, Feb. 2018 pp.82-89.
- [23] Ribaric S, Ariyaeeinia A, Pavesic N. De-identification for privacy protection in multimedia content: A survey. *Signal Processing: Image Communication*, 2016, 47: 131-151.
- [24] Gross R, Sweeney L, Cohn J, de la Torre F, Baker S. Face de-identification. In *Protecting Privacy in Video Surveillance*, Senior A (ed.), Springer, 2009, pp.129-146.
- [25] Frome A, Cheung G, Abdulkader A, Zennaro M, Wu B, Bissacco A, Adam H, Neven H, Vincent L. Large-scale privacy protection in Google street view. In *Proc. the 12th International Conference on Computer Vision*, Sept. 2009, pp.2373-2380.
- [26] Sweeney L. k -anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 2002, 10(5): 557-570.
- [27] Brkic K, Sikiric I, Hrkac T, Kalafatic Z. I know that person: Generative full body and face de-identification of people. In *Proc. the 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, Jul. 2017, pp.1319-1328.
- [28] Meden B, Malli R C, Fabijan S, Ekenel H K, Struc V, Peer P. Face deidentification with generative deep neural networks. *IET Signal Processing*, 2017, 11(9): 1046-54.
- [29] Meden B, Ziga E, Struc V, Peer P. k -same-net: k -anonymity with generative deep neural networks for face deidentification. *Entropy*, 2018, 20(1): Article No. 60.
- [30] Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. the 32nd International Conference on Machine Learning*, Jul. 2015, pp.448-456.
- [31] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In *Proc. the 18th International Conference on Medical Image Computing and Computer-Assisted Intervention*, Oct. 2015 pp.234-241.
- [32] Chopra S, Hadsell R, LeCun Y. Learning a similarity metric discriminatively, with application to face verification. In *Proc. the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Jun. 2005, pp.539-546.

- [33] Sun Y, Chen Y, Wang X, Tang X. Deep learning face representation by joint identification-verification. In *Proc. the 2014 Annual Conference on Neural Information Processing Systems*, Dec. 2014, pp.1988-1996.
- [34] Schroff F, Kalenichenko D, Philbin J. FaceNet: A unified embedding for face recognition and clustering. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2015, pp.815-823.
- [35] Parkhi O M, Vedaldi A, Zisserman A. Deep face recognition. In *Proc. the 26th British Machine Vision Conference*, Sept. 2015, Article No. 6.
- [36] Wu X, He R, Sun Z, Tan T. A light CNN for deep face representation with noisy labels. *IEEE Transactions on Information Forensics and Security*, 2018, 13(11): 2884-2896.
- [37] Lin M, Chen Q, Yan S. Network in network. arXiv:1312.4400, 2013. <https://arxiv.org/pdf/1312.4400.pdf>, Sept. 2018.
- [38] Du L, Ling H. Preservative license plate de-identification for privacy protection. In *Proc. the 2011 International Conference on Document Analysis and Recognition*, Sept. 2011, pp.468-472.
- [39] Ricanek K, Tesafaye T. MORPH: A longitudinal image database of normal adult age-progression. In *Proc. the 7th International Conference on Automatic Face and Gesture Recognition*, Apr. 2006, pp.341-345.
- [40] Kingma D P, Ba J. Adam: A method for stochastic optimization. arXiv:1412.6980, 2015. <https://arxiv.org/pdf/1412.6980.pdf>, Sept. 2018.



Yifan Wu received her B.S. degree in communication engineering from China University of Geosciences, Wuhan, and M.S. degree in computer science from Temple University, Philadelphia, in 2016 and 2017, respectively. She is currently a Ph.D. student in bioengineering at University of Pennsylvania,

Philadelphia. Her current research interests are computer vision and medical image processing.



Fan Yang received his B.S. degree in electrical engineering from Anhui Science and Technology University, Bengbu, and M.S. degree in biomedical engineering from Xi'dian University, Xi'an, in 2012 and 2015, respectively. He is currently a Ph.D. student in Department of Computer and Information Sciences at Temple University, Philadelphia.

His current research interests are computer vision and machine learning.



Yong Xu received his B.S., M.S. and Ph.D. degrees in mathematics from Nanjing University, Nanjing, in 1993, 1996 and 1999, respectively. He was a post-doctoral research fellow of computer science with South China University of Technology, Guangzhou, from 1999 to 2001, where he became a faculty member and he is currently a professor with the School of Computer Science and Engineering. His current research interests include image/video classification and recognition, object recognition, action recognition, and big data analysis. Prof. Xu is a senior member of CCF, ACM, and IEEE.



Haibin Ling received his B.S. degree in mathematics and his M.S. degree in computer science from Peking University, Beijing, in 1997 and 2000, respectively, and his Ph.D. degree from the University of Maryland, College Park, in computer science in 2006.

From 2000 to 2001, he was an assistant researcher at Microsoft Research Asia, Beijing. From 2006 to 2007, he worked as a postdoctoral scientist at the University of California Los Angeles. After that, he joined Siemens Corporate Research as a research scientist. Since the fall 2008, he has been with Temple University, Philadelphia, where he is now an associate professor. His research interests include computer vision, augmented reality, medical image analysis, and human computer interaction. He received the Best Student Paper Award at the ACM Symposium on User Interface Software and Technology (UIST) in 2003, and the NSF CAREER Award in 2014. He serves as associate editors for several journals including IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI), Pattern Recognition (PR), and Computer Vision and Image Understanding (CVIU). He has also served as Area Chairs for CVPR 2014, CVPR 2016 and CVPR 2019.